

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : <b>G06T 11/00</b>		A2	(11) International Publication Number: <b>WO 98/59321</b> (43) International Publication Date: 30 December 1998 (30.12.98)
(21) International Application Number: <b>PCT/US98/13341</b> (22) International Filing Date: 25 June 1998 (25.06.98)  (30) Priority Data: 08/882,721 25 June 1997 (25.06.97) US		(81) Designated States: JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).  Published <i>Without international search report and to be republished upon receipt of that report.</i>	
(71) Applicant: HAPTEK, INC. [US/US]; 931 Pacific Avenue, Santa Cruz, CA 95060-4429 (US).  (72) Inventors: SHAW, Christopher, D.; P.O. Box 8218, Santa Cruz, CA 95060 (US). WILSON, Orion; 1635 Bridgegate Street, West Lake Village, CA 91361 (US).  (74) Agent: O'DOWD, Shawn, W.; Kenyon & Kenyon, Suite 600, 333 W. San Carlos, San Jose, CA 95110 (US).			
(54) Title: METHODS AND APPARATUSES FOR CONTROLLING TRANSFORMATION OF TWO AND THREE-DIMENSIONAL IMAGES			
(57) Abstract			
<p>Methods and apparatuses described herein automate and confer additive properties to morphs (modification of a starting graphical image to a destination graphical image). The enhanced automated additive morphs created by this invention extend the currently limited scope of animation techniques, creating: moving morphs, where characters can speak, move, and emote during the morphing process; parametric character creation, where features can be sequentially added to a character to create a wide variety of resulting characters; behavioral transference, where character behavior can be automatically transferred to newly created characters, and behavioral layering whereby sequential behavior patterns can be concurrently transferred or imparted to a character. The present invention allows an animator to create, animate, control and transform two- and three-dimensional images instantaneously and fluidly. The invention provides a superior solution at significantly less cost which extends the range and properties of existing state of the art animation.</p>			
<pre>graph TD     931[931] --&gt; 932[932]     932 --&gt; INPUT     INPUT --&gt; 937[937]     937 --&gt; 938[938]     938 --&gt; 940[940]     940 --&gt; 939[939]     939 --&gt; 933(( ))     939 --&gt; 934[934]     934 --&gt; 935(( ))     935 --&gt; OUT[TO SCREEN, VIDEO, FILE EXPORT, ETC.]     933 --&gt; OUT     933 --&gt; OTHER[OTHER SYSTEMS, IE: WWW]</pre>			

REFT AVAIL ARIF COPY

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

## METHODS AND APPARATUSES FOR CONTROLLING TRANSFORMATION OF TWO AND THREE-DIMENSIONAL IMAGES

### Field of the Invention

The present invention pertains to automated methods and apparatuses for the controlling and transforming of two and three-dimensional images. More particularly, the present invention relates to methods and apparatuses for changing the elements of image 5 through the use of one or more sets of modification data in real time.

### Copyright Notice

A portion of the disclosure of this patent document contains material which is subject 10 to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or patent disclosure as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

### Background of the Invention

15 Referring to Fig. 8, a computer system that is known in the art is shown. The computer system 810 includes a system unit having a processor 811, such as a Pentium®

processor manufactured by Intel Corporation, Santa Clara, California. The processor is coupled to system memory 812 (e.g., Random Access Memory (RAM)) via a bridge circuit 813. The bridge circuit 813 couples the processor 811 and system memory 812 to a bus 814, such as one operated according to the Peripheral Component Interconnect standard (Version 2.1, 1995, PCI Special Interest Group, Portland, Oregon). The system unit 810 also includes a graphics adapter 815 coupled to the bus 814 which converts data signals from the bus into information for output at a display 820, such as a cathode ray tube (CRT) display, active matrix display, etc. Using the computer system of Fig. 1, a graphical image can be displayed at display 820. The graphical image can be created internally to the computer system 810 or can be input via an input device 830 (such as a scanner, video camera, digital camera, etc.). As is known in the art, a graphical image is stored as a number of two-dimensional picture elements or "pixels," each of which can be displayed.

In the current art, graphical images (e.g., of a person's face) can be changed by allowing the user to modify a graphical image by "moving" (e.g., with a cursor movement device such as a mouse) the two-dimensional location of one or more pixels (For example: 15 Adobe Photoshop Version 3.0.5 (Adobe Systems, Inc., San Jose, California)). In doing so, the other pixels around the one that is being moved are filled in with new data or other pixel data from the graphical image. For example, the graphical image of the person's face can be modified using this product by making the person's nose larger or smaller. This two-dimensional phenomenon is analogous to stretching and warping a photograph printed on a 20 "rubber sheet." In the Kai's Power Goo product by MetaTools, Inc (Carpinteria, California), photographic distortions can be performed in "real time" by the operator's "clicking and dragging" with a mouse across the surface of a photo displayed on the computer screen. The operator can see the photograph stretch as the mouse is moved. This procedure covers only 25 two dimensional art and does not permit any sophisticated character animation such as speech or emotion.

In the current art, the gradual change of the shape of one image into that of another as seen in film and video is called a "morph". Current morphs are created by an operator who instructs a computer to distort the shape of a specific starting image into the shape of a 30 specific target image. Morphing programs typically work by allowing the operator to select points on the outline of the specific starting image and then to reassign each of these points to

a new location, thereby defining the new outline of the desired target image. The computer then performs the morph by: (1) smoothly moving each of these points along a path from start to finish and (2) interpolating the movement of all the other points within the image as the morph takes place.

5        There are two distinct disadvantages to this method described above. First, it requires that a custom morph be created for each desired transformation. Second, because this method requires the selection of a single image or frame upon which the morph is performed, the frame-by-frame progression of character action must stop during the period in which the morph is performed. This is why in current films, characters do not speak or move during the  
10      morph procedure. The reason morphs are currently performed relatively quickly (i.e., within a few seconds) is so that this freezing of action is not fully noticed by the audience.

In recent films, whenever a character morphs (e.g. when the villain robot in James Cameron's "Terminator 2" changes to its liquid metal form), the character ceases moving while the morph takes place. In the "Fifth Element" released in May of 1997, characters are  
15      seen changing from alien to human form while they shake their heads back and forth. Although this gives the character the appearance of moving while the morph is taking place, the underlying 3D image of a character's head is actually frozen while it shakes. This method is merely the 3D equivalent of a "freeze frame." This method cannot enable a morphing character to speak, move or emote while a morph is taking place. The static morphing  
20      methods used in today's films are slow and considerably expensive.

#### Summary of the Invention

According to an embodiment of the present invention a first region of a first graphical image is identified and then it is modified based on a first set of predetermined modification data. Using this method to morph a graphical image, a variety of applications can be  
25      performed according to further embodiments of the present invention.

First, the morph (e.g., the application of modification data) for a first starting image can be readily applied to other starting images. In other words, the morphs automatically impart desired characteristics in a custom manner to a multiplicity of starting images. This an  
30      improvement over the prior art which requires a customized selection and reassignment of points on a specific starting image to create a morph. A method of the present invention

described herein automates this process. Rather than requiring an artist or technician to custom create a morph of a specific image, for example, an adult into that of a child, the morphs of the present invention enable a wide variety of human, animal, or other characters to be rendered chimplike using a single "chimp" morph. An example of this is shown in Fig. 1, where a single morph relating to "chimpification" is applied to three different starting images. In each case, the resulting image maintains recognizable features of the starting image while uniquely embodying the desired characteristics of the "chimp" morph. The morph has been described thus far as modification data. Examples of modification data includes deltaset and deltazones described in more detail below. Briefly, deltaset or zones categorically identify regions, feature by feature within differing starting images so that these images are uniquely altered to preserve the automated morph's desired effect. Because a single morph can be applied to a number of different starting images, the morph exists as a qualitative entity independently from the images it acts upon. This independence creates an entirely new tool, a **morph library**, a collection of desired alterations or enhancements which can be generically used on any starting image to create specific desired effects as illustrated in the above "chimpification" example.

Second, once an image has been morphed to add a particular characteristic or quality, the resulting image can be subjected to a different morph to add a second characteristic or quality. Fig. 2 illustrates a simple example of this additive property wherein a "chimp" morph is added to a "child" morph to create a childlike chimp (other examples will be described in further detail below). The additive property of the automated, additive morphing system can be used in a number of ways to bring new functionality and scope to the morphing process. Five distinct additive properties of automated, additive morphs will be described below along with their practical application.

Third, morphs can be provided that allow a graphical image character to speak, move, emote, etc. According to an embodiment of the invention, a moving morph can be created during which a character can continue speaking, moving, and emoting by cross applying an automated additive morph to a ("morph sequence"). The morph sequence that is known in the art (such as what is shown in programs by Dr. Michael Cohen at the University of California at Santa Cruz and products of Protozoa Inc. (San Francisco, California) allows for computer-generated characters to move their mouths in a manner which approximates speech

by running their characters through a sequence of "Viseme" morphs. (A Viseme is the visual equivalent of a phoneme, i.e. the face one makes when making a phonetic sound.) Such programs use a specific initial image of a character at rest, and a collection of target images. Each target image corresponds to a particular facial position or "Viseme" used in common speech. Fig. 3 shows how these target images can be strung together in a morph sequence to make an animated character approximate the lip movements of speech. This figure shows the sequence of Viseme endpoints which enable the character to mouth the word "pony". It is important to note that the creation and performance of this sequence does not require the special properties of the morphing system presented herein. The morphs within this sequence in the prior art are not generalized (all are distortions of one specific starting image), and they are not additive. Visemes used in the prior art follow one after the other and are not added to one another. According to an embodiment of the present invention, the use of morph sequences is extended to generate not only speech, but also emotive flow and the physical components of emotional reactions.

For example, Fig. 4 shows the cross-addition of an automated, additive morph to the "pony" morph sequence described above in Fig. 3. In that figure, the four vertical columns of character pictures represent the progressive application of the "chimp" morph described earlier (from left to right multiplying the modification data by multiplication values of 0%, 33%, 66%, 100% prior to application to the starting image). Because the "chimp" morph is non-specific as to its starting point (as are all automated additive morphs according to the present invention) it is possible to increasingly apply the "chimp" morph while changing the starting point within the morph Sequence, producing the progression shown in the darkened diagonal of squares. This diagonal progression, shown in horizontal fashion at the bottom of Fig. 4 yields a character which can speak while this character is morphing. This is the underlying structure of the moving morph. Traditional morphs, (being specific rather than generic) cannot be cross-applied in this manner. Characters created using the methods of the present invention can be made to not only speak, but also emote, and react from morph sequences. Thus, characters can remain fully functional during an automated, additive morph rather than being required to "freeze frame" until the morph has been completed as do the morphs of the prior art. An additional benefit of this cross-additive procedure is that morphs can be stopped at any point to yield a fully functional, consistent new character which is a

hybrid of the starting and final characters.

Fourth, the methods of the present invention provide for parametric character creation in which newly created characters automatically speak, move, and emote using modification data stored in a database or library. In the automated, additive morphing system, (1) morphs 5 can exist as qualitative attributes independent of any particular starting image, and (2) morphs can be applied, one after the other, to produce a cumulative effect. When qualitative attributes are defined as appearance parameters (length or width of nose, prominence of jaw parameters, roundness of face, etc.) these attributes can be selectively applied in such a way as to create any desired face from one single starting image. As an example, a multi-racial starting 10 character is defined and a morph library of appearance parameters is created which can be used to adjust the characters features so as to create any desired character. **Fig. 5** shows an example of this process. The parameter adjustments in this figure are coarse and cartoon-like so as to yield clearly visible variations. In realistic character generation, a much larger number of parameters can be more gradually applied. The first three morphs shown in this 15 illustration are "shape morphs." In the final step of Figure 5, the coloration or "skin" which is laid over the facial shape is changed rather than the shape itself. This step can be used to create the desired hair color, eye color, skin tone, facial hair, etc. in the resultant character.

Fifth, the parametric character creation described above can be combined with the moving morph, also described above, to create characters which automatically speak, emote 20 and move. This dual application is illustrated in **Fig. 6**, wherein not only the underlying structure, but also the full speaking and emoting functionality of the original character are automatically transferred to the new character. The character shown in Fig. 6 not only contains a mutable physical appearance, but also a full set of Visemes, emotions, and computer triggered autonomous and reactive behavior. All of these functions can be 25 automatically transferred to a vast range of characters which can be created using parametric character creation. This represents an exponential savings in animation time and cost over existing procedures which require custom creation of not only the character itself, but every emotion, Viseme, blink, etc. that the new character makes.

Sixth, morph sequences can be used simultaneously to combine different behavioral 30 sequences. **Fig. 7** illustrates the simultaneous utilization of an emoting sequence and a speaking sequence. In this figure, the Viseme sequence required to make the character say

"pony" (left column of pictures) has been added to an emotive morph sequence (center column) in such a manner that the timing of each sequence is preserved. The resultant sequence (right column) creates a character which can simultaneously speak and react with emotions. This procedure can also be used to combine autonomous emotive factors (a computer-generated cycling of deltas representing different emotions or "moods") with reactive factors (emotional deltas triggered by the proximity of elements within the character's environment which have assigned emotive influences on the character). Such procedures can be used to visualize the interplay between conscious and subconscious emotions.

The foregoing examples and other examples of the present invention will be described in further detail below.

#### Brief Description of the Drawings

**Fig. 1** shows an example of an automated morph according to an embodiment of the present invention.

**Fig. 2** shows an example of an automated, additive morph according to an embodiment of the present invention.

**Fig. 3** shows an example of a morph sequence that can be performed according to an embodiment of the present invention.

**Fig. 4** shows an example of a moving morph according to an embodiment of the present invention.

**Fig 5** shows an example of parametric character creation according to an embodiment of the present invention.

**Fig. 6** shows an example of automatic behavioral transference according to an embodiment of the present invention.

**Fig. 7** shows an example of behavioral layering according to an embodiment of the present invention.

**Fig. 8** is a computer system that is known in the art.

**Fig. 9** is a general block diagram of an image transformation system of the present invention.

**Figs. 10 a-d** are polygonal models used for the presentation of a graphical image of a

human head or the like.

**Figs. 11 a-f** are polygonal images showing the application of delset in accordance with an embodiment of the present invention.

**Figs. 12 a-g** are graphical images of a person's head that are generated in accordance with an embodiment of the present invention.

**Fig. 13** shows an input device for controlling the amount of transformation occurs when applying a delset to an image.

**Figs. 14 a-d** are graphical images of a person's head that are generated in accordance with an embodiment of the present invention.

**Fig. 15** shows a communication system environment for an exemplary method of the present invention.

#### Detailed Description

According to an embodiment of the present invention modification data is generated that can be applied to a starting image so as to form a destination image. For example, the modification data can be difference values that are generated by determining the differences between first and second images. Once these differences are determined they can be stored and later applied to any starting image to create a new destination image without the extensive frame-by-frame steps described above with respect to morphing performed in the motion picture industry. These difference values can be created on a vertex-by- vertex basis to facilitate the morphing between shapes that have an identical number of vertices.

Alternatively, difference values can be assigned spatially, so that the location of points within the starting image determines the motion within the automated morph. This eliminates the need for explicit identification of vertices and allows these methods and apparatuses to work regardless of a given image's polygonal structure. For simplicity sake we describe an example vertex-based automated additive morphing system below which uses delsets as the modification data. A position or spatially-based morphing system which can morph images regardless of polygonal structure, such as delzones (another example of modification data), is created by interpolating the motion between vertices.

An example of the vertices-based embodiment of the present invention includes the generation of a first image (e.g., a neutral or starting image) comprising a first number of

vertices, each vertex having a spatial location (e.g., in two or three dimensional space) and a second image is generated (e.g. a target or destination image) having an equal number of vertices. A difference between a first one of the vertices of the first image and a corresponding vertex of the second image is determined representing the difference in  
5 location between the two vertices. The difference is then stored in a memory device (e.g., RAM, hard-disc drive, etc.). Difference values for all corresponding vertices of the first and second images can be created using these steps and stored as a variable array (referred to herein as a deltabase). The deltabase can then be applied to the first image to create the second image by moving the vertices in the first image to their corresponding locations in the second  
10 image. Alternatively, a multiplication or ratio value can be multiplied by the entries in the deltabase and applied to the first image so that an intermediate graphical image is created.  
According to a feature of the present invention, the deltabase can be applied to any starting image having an equal number of vertices. This allows the user to create new destination images without performing, again, the mathematical calculations used to create the original  
15 deltabase.

Referring to Fig. 9, a general block diagram of an image transformation system 900 of the present invention is shown. According to an embodiment of the present invention, the system 900 includes a library or database of deltabases 931. The library of deltabases 931 can be stored in the system memory 912 or any other memory device, such as a hard disc drive 917  
20 coupled to bus 914 via a Small Computer Standard Interface (SCSI) host bus adapter 918 (see Fig. 8). As described in further detail below, deltabases are variable arrays of position change values that can be applied to the vertices of a starting image. Referring to Fig. 9, the deltabase information is composed and cached in device 932 (e.g., processor 811 and system memory 812 of Fig. 8) where it then can be used to transform a first or starting image having a neutral  
25 geometry 933 into a target image having a final geometry 935. Additional geometry manipulation can be performed such as the addition of features (e.g., hair) or actions (e.g., looking around) by device 934. Both the starting and second images can then be displayed at display 920, or any other output device (memory) or sent to file export (e.g., the Internet system).

30 Inputs to the system 900 of Fig. 9 include a variety of user controls 937; autonomous behavior control 938; and face tracker data input 939 which will be further described below.

Other inputs can come from other systems such as the so-called World-Wide Web (WWW). Also audio data can be supplied by audio data input device 940 which can be supplied to deltaset caching and composing device 932.

In this embodiment, the neutral geometry 933 is based on the image of a person's head that has been captured using any of a variety of known methods (e.g., video, scanner, etc.). Referring to Fig. 10, the image data of the person's head is placed onto a polygonal model 1051. The polygonal model comprises a plurality of vertices 1052 and connections 1053 that extend between the vertices. Each polygon 1054 of the polygonal model is defined by three or more vertices 1052. To show the generation and application of deltaset to the polygonal model of Fig. 10a, an example is discussed below using simple polygons (e.g., a square, a triangle, a rectangle, and a circle).

Each polygon has an identifiable shape. For example, looking at Fig. 11a, a square polygon is shown having 8 vertices (points 1100 to 1107) in two-dimensional space. By moving individual vertices, the square polygon can be converted into a number of other polygon shapes such as a rectangle (Fig. 11b), a circle (Fig. 11c) and a triangle (Fig. 11d; where vertices 1100, 1101, and 1107 all occupy the same point in two-dimensional space). A deltaset is a set of steps that are taken to move each vertex (1100 to 1107) from a starting polygon to a target or destination polygon. For example, the steps that are taken from the square polygon of Fig. 11a to the rectangular polygon of Fig. 11b include vertices 1105, 1106, and 1107 moving to the left a certain distance "x"; points 1101, 1102, and 1103 moving to the right the same distance "x"; and vertices 1100 and 1104 staying in the same location. Thus, the deltaset defines the path taken by each vertex in transforming the starting polygon to the destination polygon. In other words, the deltaset defines the difference in position of corresponding vertices in the starting and target polygons. Similarly, deltaset can be created for the transformation of the square polygon of Fig. 11a to the circle polygon of Fig. 11c and of the square polygon of Fig. 11a to the triangle polygon of Fig. 11d.

In this embodiment, the deltaset is created by transforming a starting polygon shape into another, however, one skilled in the art will appreciate that a deltaset can be created that are not based on specific starting and target shapes, but created in the abstract. Moreover, once a deltaset is created, it can be used on any starting shape to create a new shape. For example, the deltaset used to transform the square polygon of Fig. 11a to the rectangle

polygon of Fig. 11b (for convenience, referred to as Deltaset1) can be used on the circle polygon of Fig. 11c. Thus, the circle polygon of Fig. 11c becomes the starting shape and after applying Deltaset1, would become the ellipse polygon of Fig. 11e (i.e., the target shape).

Deltasets can also be combined (e.g., added together) to create new deltaset. Thus,  
 5 Deltaset1, Deltaset2 (i.e., transform from the square of Fig. 11a to the circle of Fig. 11c), and  
 Deltaset3 (i.e., transform from the square of Fig. 11a to the triangle of Fig. 11c) can be  
 combined to form a new deltaset (Deltaset4). Applying Deltaset4 to the starting square  
 polygon of Fig. 11a, the target shape of Fig. 11f is achieved. In its simplest form, the starting  
 polygon, destination polygon, and deltaset must have the same number of vertices. Additional  
 10 algorithms would be necessary to transform between shapes or objects having a differing  
 number of vertices.

An additional method for moving vertices can be derived from the deltaset method  
 wherein the motion to the points of a deltaset are interpolated such that a continuous field of  
 motion is created. These fields which we refer to as deltazones can be used to morph images  
 15 irrespective of their particular triangle strip set because a one to one correspondence between  
 movements and vertices upon which the deltaset rely are replaced by a dynamical system of  
 motion which operates on any number of vertices by moving them in accordance with their  
 original location.

Herein, an example of the implementation of deltaset and their operation on  
 20 graphical images will be described with reference to pseudocode based on "C" and "C++"  
 programming that is known in the art. The datatype structure for a deltaset (Deltaset\_Type) is  
 similar to that for a basic shape object, and the pseudocode is shown in Table I.

```
//////////  

// Basic datatype structure of a deltaset.  

25  typedef struct {  

        dataPoint_Type *dataPoints; // Array of delta values  

        int             numPoints; // Number of points in above  

    } deltaSet_Type, shape_Type;  

// end of datatype structure  

30  //////////
```

Table I

As seen from the above, the Deltaset\_Type and shape\_Type variables each include an  
 35 array of [numpoints] values. Each value is a position of a vertex for the shape\_Type variable

and delta value for the Deltaset\_Type variable.

An example of a core routine for the creation of a deltaset through the calculation of the steps or difference values from a starting object to a destination object is shown in Table II.

```
5 ///////////////////////////////////////////////////////////////////
// core Routine to calculate the steps from a source(neutral)
// object to a destination object and store those steps
// in a deltaset.

10 DeltaSet_Calc(deltaSet_Type    *dset,
                 shape_Type      *src
                 shape_Type      *dest)

15 {
    int             i;
    int             numpts;

20     dataPoint_Type delta;

    // Ensure that dset has a matching number of data points
    // as the shapes.
    Numpts = src -> numPoints;
    deltaSet_SetNumPts(dset, numpts);

    // For each data point in the objects, calculate the
    // difference between the source and the destination
    // and store the result in the deltaset.

25     for (i = 0; i < numpts; i++)
    {
        delta = dest -> dataPoints[i] - src -> dataPoints[i];
        dset -> dataPoints[i] = delta;
    }
30 }

// end of routine
```

```
//////////
```

Table II

5 As can be seen from the above pseudocode, the variable "delta" is used to temporarily store the difference in position between the source (src) and destination (dest) for each of the vertices in the shape. Each delta value is then stored in a deltaset array (dset). Once a deltaset array is created, it can be easily applied to any starting shape having an equal number of vertices to form a new target shape.

10 An example of pseudocode that can be used to modify a deltaset so that it can be applied to a starting shape is shown in Table III.

```
//////////
// core Routine to apply the steps stored in a
15 // deltaset to a shape., with a percentage amount.
// Note that negative amounts can be used.
deltaSet_Apply (deltaSet_Type *dset,
                 shape_Type      *dest,
                           float           amount)
20 {
    int   i;

    if (amount == 0.0)
        return;
25
    for (i = 0 ; i < dset -> numPoints; i++)
    {
        dest -> dataPoints[i] += (dset -> dataPoints [i] *
30     amount);
        }
    }
    // end of routine
//////////
```

35

Table III

As seen from above, during the routine Deltaset\_Apply, calculations for a single transition are performed based on a percentage amount passed using the "amount." Each data

point in the destination shape is calculated based on the deltaset value for that point multiplied by the percentage value "amount" (which can have a negative value or a value greater than 1).

5       The pseudocode of Table IV shows two utility routines that are used for creating a new, blank deltaset and to set the number of datapoints.

```
//////////  
// utility Routine to create a new, Blank deltaset  
deltaSet_Type *NewDeltaSet()  
10    {  
    allocate a new deltraSet_Type and return a pointer to it.  
  }  
  // end of routine  
//////////  
15  
//////////  
// utility Routine to set the number of datapoints  
// in a deltaset.  
deltaSet_SetNumPts (deltaSet_Type *dset, int numPoints)  
20    {  
    de-allocate dset -> dataPoints, if not already empty;  
  
    allocate an array of type dataPoint_Type and size numPoints,  
    and put it in dset -> dataPoints;  
25  
    dset -> numPoints = numPoints;  
  }  
  // end of routine  
//////////  
30
```

Table IV

With reference to Table V, an example of pseudocode is shown for the transformation  
 5 from the square shape of Fig. 11a to the shape of Fig. 11e.

```

////////// The following pseudocode example shows how
// to use the above deltaset routines to morph from a
// square to a new shape which has features of both the
10 // rectangle and the circle.

// Define the basic dataPoint that makes shapes & deltaset.
typedef 2DVector dataPoint_Type;

15 // Declaration of basic shapes.
shape_Type      Square,
                Rectangle,
                Circle;

20 // Declaration of deltaset.
deltaSet_Type   rect_dset,
                circ_dset;

25 // Declaration of a new shape to get the data put into it.
shape_Type      newshape;

// Initialize shapes.

30 shape_SetNumPoints(&square,          8);
shape_SetNumPoints(&rectangle,        8);
shape_SetNumPoints(&circle,           8);

// set data points of square, rectangle, and circle.

35 shape_SetPoints (&square,      ( 0.0,    1.0),
                  ( 1.0,    1.0),
                  ( 1.0,    0.0),
                  ( 1.0,   -1.0),
                  ( 0.0,   -1.0),
40                  (-1.0,   -1.0),
                  (-1.0,    0.0),
                  (-1.0,    1.0));

shape_SetPoints (&rectangle,   ( 0.0,    1.0),

```

```
5           ( 2.0,    1.0),
           ( 2.0,    0.0),
           ( 2.0,   -1.0),
           ( 0.0,   -1.0),
           (-2.0,   -1.0),
10          (-2.0,    0.0),
           (-2.0,    1.0));
shape_SetPoints (&circle,      ( 0.0,    1.0),
15          ( 0.0,    1.0),
           ( 0.0,    1.0),
           ( 0.0,    1.0),
           ( 0.0,    1.0),
           ( 0.0,    1.0),
           ( 0.0,    1.0),
           ( 0.0,    1.0));
20 // calculate DeltaSets
deltaSet_Calc (&rect_dset, &square, &rectangle);
deltaSet_Calc (&circ_dset, &square, &circle);

///////////////////////////////
25 // The resulting DeltaSets now contain the values:
// rect_dset:  ( 0.0,    0.0)
//              ( 1.0,    0.0)
//              ( 1.0,    0.0)
//              ( 1.0,    0.0)
30 //              ( 0.0,    0.0)
//              (-1.0,    0.0)
//              (-1.0,    0.0)
//              (-1.0,    0.0)
//
35 //circ_dset:   ( 0.0,    0.0)
//              (-0.4,   -0.4)
//              ( 0.0,    0.0)
//              (-0.4,    0.4)
//              ( 0.0,    0.0)
40 //              ( 0.4,    0.4)
//              ( 0.0,    0.0)
//              ( 0.0,   -0.4)
///////////////////////////////
45 // Apply the DeltaSets.
newshape = copy of square;
```

```

deltaSet_Apply(&rect_dset, &newshape, 1.0);
deltaSet_Apply(&circ_dset, &newshape, 1.0);

5   ///////////////////////////////////////////////////////////////////
// newshape now contains values which
// look like the ellipse drawn above:
//      ( 0.0,    1.0)
//      ( 1.6,    0.6)
10  //
//      ( 2.0,    0.0)
//      ( 1.6,   -0.6)
//      ( 0.0,   -1.0)
//      (-1.6,   -0.6)
//      (-2.0,    0.0)
15  //
//      (-1.6,    0.6)
///////////////////////////////////////////////////////////////////

// to create the egg-ish shape above,
// one would simply add a third DeltaSet
20  // based on a triangle shape.

-----end of simple geometry example///////////

```

Table V

As seen from the pseudocode of Table V, a datapoint is defined as a two-dimensional vector and the square, rectangle, and circle shapes are defined as eight points with abscissa and ordinate values. Deltasets are then calculated for the transition from square to rectangle and from square to circle. As seen above, the resulting deltasets (rect\_dset and circ\_dset) represent differences between abscissa and ordinate values of the respective starting and target images. The deltasets can then be applied to a starting shape (in this example, the starting image, newshape, is set to the square shape of Fig. 11a). First the rect\_dset deltaset is applied to the square shape to form an intermediate shape, and then the circ\_dset deltaset is applied to this intermediate shape to form the destination shape that is shown in Fig. 11e. To get to the shape of Fig. 11f, a deltaset representing a transformation between the square shape of Fig. 11a to the triangle shape of Fig. 11d is created and applied to the ellipse shape shown in Fig. 11e.

The deltasets example, above, can be easily extended to a three-dimensional representation. The example can also be expanded to more intricate and complex applications

such as in three-dimensional space and facial animation. As the application of deltaset becomes more complex in facial animation, several additional features can be added. For example, certain motions of the face are limited to certain defined areas, such as blinking of the eyes. Accordingly, a deltaset for an entire face would be mostly 0's (indicating no change) except for the eyes and eyelids, thus isolating these areas for change. To improve efficiency, the deltaset datatype can be changed so that only non-zero values are stored. Thus during the execution of the Deltaset\_apply routine, only the points that change are acted upon, rather than every point in the graphical representation.

An embodiment of facial animation is described below with reference to the  
10 pseudocode example of Table VI.

Facial Moving-Morphing example -----

```
// The following pseudocode example shows how deltaset
// morphing is used to fully animate and morph a face.
15 // Note that this achieves a "moving morph", wherein the
// overall structure of the face can smoothly change without
// interrupting the process of other facial animation such
// as blinking, emoting, and speaking.

20 /////////////////
// Setup

25 // Define the basic dataPoint that makes shapes & datasets.
typedef 3DVector dataPoint_Type;

30 // Declaration of basic shapes & deltatasets.
shape_Type      neutralFace,
                overallMorphface,
                blinkFace,
                emoteFaces [],
                speakFaces [],
                newShapeFace;

35 deltaSet_Type overall_dset,
                blink_dset,
                emote_dsets[],
                speak_dsets[];

40 // neutralFace is the geometry of the basic 3D face,
//   no expression, looking straight ahead.
// overallMorphFace is a radically different face,
```

```
// say a cat.  
// blinkFace is the same as neutralFace but with eyes closed.  
// emoteFaces is an array of faces with different emotions or  
// expressions. ie Happy, sad, angry, trustful etc.  
5 // speakFaces is an array of faces in different phoneme (or  
// "viseme") positions. ie "OO", "AE", "L", "M" etc.  
// newShapeFace is a shape which is the destination of the  
// morphing.  
  
10 // Declarations of amount of morphs.  
// These typically range from 0.0 to 1.0, but can be outside  
// of this range.  
float overallMorphAmount, blinkAmount,  
15     emoteAmounts[], speakAmounts[];  
  
// Other declarations  
float time;           // a pseudo time variable.  
int numEmotes;        // the number of emotion faces.  
20 int numSpeaks;       // the number of viseme faces.  
  
/////////////////////////////  
// initialize the deltaset  
25 deltaSet_Calc(&overall_dset, &neutralFace, &blinkFace);  
deltaSet_Calc(&blink_dset, &neutralFace, &blinkFace);  
for (i = 0; i < numEmotes; i++)  
    deltaSet_Calc( &emote_dsets[i],  
                  &neutralFace,  
30                  &emoteFaces[i]);  
for (i = 0; i < numSpeaks; i++)  
    deltaSet_Calc( &speak_dsets[i],  
                  &neutralFace,  
                  &speakFaces[i]);  
35 ///////////////////  
// Main animation loop  
while (KeepRunning)  
{  
40     time += 0.1;  
  
    // Calculate the amount each morph is to be applied.  
    // For emoteAmounts and speakAmounts, this is an array  
    // of values most of which are zero.  
45    // (Note that deltaSet_Apply() returns immediately if  
    // amount == 0.0)
```

```
5      CalBlinkAmount (&blinkAmount);
CalEmoteAmounts (emoteAmounts);
CalcSpeakAmounts (speakAmount);
overallMorphAmount = sin(time) * 0.5 + 0.5;
10     // Reset the working copy of the face.
newShapeFace = Copy of neutralFace;
// Apply the data sets controlling facial animation.
deltaSet_Apply(blink_dset, &newShapeFace, blinkAmount);
for (i = 0; i < numEmotes; i++)
    deltaSet_Apply(&emote_dsets[i],
                    &newShapeFace,
                    emoteAmounts [i]);
15     for (i = 0 i < numSpeaks; i++)
        deltaSet_Apply(&speak_dsets[i],
                        &newShapeFace,
                        speakAmounts [i]);
20     // Apply the overall shape morph
deltaSet_Apply(&overall_dset,
                &newShapeFace,
                overallMorphAmount);
}
// End of animation loop
///////////////////////////////
25 //-----end of facial moving morph example-----
```

Table VI

30 As seen from the above, the animated face image comprises three-dimensional  
datapoints. “NeutralFace” is a starting image that will be changed based on one or more  
deltasets. The neutralface image is shown in Fig. 12a with eyes looking straight ahead and no  
expression. “OverallMorphFace” is a different face from NeutralFace. In this example,  
35 OverallMorphFace is in the image of a cat shown in Fig. 12b. A face showing a completed  
facial movement is “blinkFace” which shows the same face as NeutralFace but with the eyes  
closed (see Fig. 12c). “EmoteFaces” is an array of the neutralFace augmented to show one or  
more emotions. For example, Fig. 12d shows the neutralFace emoting happiness, Fig. 12e  
shows neutralFace emoting anger, etc. “SpeakFaces” is an array of faces showing expressions  
40 of different phonemes. a phoneme, or viseme, is a speech syllable used to form spoken words  
(e.g., the “oo”, “ae”, “l”, and “m” sounds). As an example, Fig. 12f shows neutralFace

expressing the phoneme "oo."

The amount of transformation or morphing can be controlled by multiplication or multiplier values, overallMorphAmount, blinkAmount, emote Amounts[], and speakAmounts[]. As an example, if blinkAmount is set to 1.0 then when applying a delset for blinking to neutralFace of Fig. 12a will achieve the face of Fig. 12c (i.e., 100% of the blink is applied). Numbers less than or greater than 1.0 can be selected for these variables.

Datasets are then created for transforming the neutralFace image. As can be seen from the pseudocode of Table VI, delset overall\_dset is created for the changes between neutralFace (Fig. 12a) and overallMorphFace (Fig. 12b); delset blink\_dset is created for the changes between neutralFace (Fig. 12a) and blinkFace (Fig. 12c); delsets emote\_dsets[] are created between neutralFace (Fig. 12a) and each emotion expression image (e.g., the "happy" emoteFace[] of Fig. 12d and the "angry" emoteFace[] of Fig. 12e; and delsets speak\_dsets[] are created between neutralFace (Fig. 12a) and each phoneme expression image (e.g., the "oo" speakFace[] of Fig. 12f).

In the main animation loop, the amounts for each delset transformation are calculated (e.g., the values for overallMorphAmount, blinkAmount, emoteAmount[], and speakAmounts[]). For the emoteAmounts[] and speakAmounts[] arrays, these values are mostly zero. The new facial image to be created is stored in newShapeFace and is originally set to the NeutralFace image. Then, the delsets that were calculated above, are applied to the newShapeFace in amounts set in transformation variables calculated above. In this example, overallMorphAmount is set to 0.5 (i.e., halfway between neutralFace and overallMorphFace; blinkAmount is set to 1.0 (i.e., full blink - eyes closed); emoteAmount[] for "happy" is set to 1.0 while all other emoteAmount[] values are set to 0; and speakAmount[] for the phoneme "oo" is set to 1.0 while all other speakAmount[] values are set to 0. The resulting image based on these variables is shown in Fig. 12g.

As described above, the delsets that have been created can now be applied to another starting image (i.e., an image other than neutralFace shown in Fig. 12a) without recalculation. This is shown in the examples of Figs. 14a-d. Using the method set forth above, a delset can be created between neutralFace and overallMorphFace which signifies the changes between a male human face (shown in Fig. 12a) and the face of a cat (shown in Fig. 12b). As seen in Fig. 14a, a neutral, male human face is shown without application of this delset. Fig. 14b

shows the effects of the application of this deltaset (or fractional value of this deltaset) in that the male human face now looks "cat-like." The underlying polygonal model for Figs. 14a and 14b are shown in Figs. 10a and 10b, respectively. As seen in Figs. 10a and b, vertices of the first image are shown to move to different positions in the destination image. Referring back 5 to Figs. 14a and 14b, one skilled in the art will appreciate that the color of each pixel can also change in accordance with a deltaset storing the difference in color for each pixel in the human and cat images of these figures. The deltaset described above can be applied to a neutral, female human face (see Fig. 14c) to form a new destination image (see Fig. 14d).

Also, the variables (e.g., overallMorphAmount) can be input using graphical sliders 10 shown in Figs. 13a-d. In this example, several deltatasets have been previously determined. a first deltaset represents the difference between a starting image with lips in a first position and a target image with lips in a second, higher position. a second deltaset represents the difference between a starting image with jaw in a first position and a target image with the jaw in a second, jutted-out position. a third deltaset represents the difference between a 15 starting image with relatively smooth skin and a target image with old (i.e., heavily textured skin). Referring to Figs. 13a-d, the amount each of these first, second, and third deltatasets is applied to the neutral image of Fig. 13a is determined by the placement of one or more sliders 1301-1303. In this example, if the slider is in a central position, then the deltaset is not applied at all (i.e., the deltaset multiplied by 0.0 is applied to the image). If the slider is placed 20 to the right, the deltaset multiplied by 1.0 is applied to the image and if it is placed to the left, the deltaset multiplied by -1.0 is applied to the image. Accordingly in Fig. 10a, sliders 101-03 are in a central position. In Fig. 13b, slider 1301 is moved (e.g., with a mouse, not shown) to the right causing the first deltaset (multiplied by 1.0) to be applied to the neutral image of Fig. 13a (thus, the lips are moved up some distance). Likewise, in Fig. 13c, slider 1302 is moved 25 to the left, and the second deltaset described above (multiplied by -1.0) is applied to the image of Fig. 13b (thus, the jaw is recessed). Also, in Fig. 13d, slider 1303 is moved to the right causing the third deltaset (multiplied by 1.0) to be applied to the image of Fig. 13c. One skilled in the art will appreciate that the sliders 1301-03 can have intermediate values between -1.0 and 1.0 or can have values beyond this range.

30 As seen from the above, once one or more deltatasets have been created, the multiplier values that are controlled by sliders 1301-03 (for example) of the embodiment of Fig. 13

would be the only input necessary to modify a starting image. This feature is advantageous in the area of communications. Referring to Fig. 15, a communication system is shown. In this system, a first component (such as server 1510) is coupled via a transmission medium 1509 to a second component (such as client 1511 coupled to a display 1512). In this example; the 5 transmission medium 1509 is the so-called Internet system that has a varying, but limited bandwidth. The server 1510 and client 1511 are computer systems similar to system 801 of Fig. 8. a first image (e.g., a person=s face) is transmitted over the transmission medium 1509 from the server 1510 to the client as well as any desired deltaset (as described above). Some code may also be sent, operating as described herein. The image and deltaset can be stored at 10 the client 1511 and the image can be displayed at display 1512. For the server 1510 to change the image at the client 1511, an entire, new image need not be sent. Rather, the multiplication values for the deltaset (e.g., the values controlled by sliders 1301-03 in Fig. 13) can be sent over the transmission medium 1509 to cause the desired change to the image at display 1512. Thus, a great savings in bandwidth is achieved allowing greater animation and control of the 15 image.

In another example, the system of Fig. 15 can be used as a video phone system where the original image that is sent is that of the speaking party at the server 1510 over the transmission medium 1509 (e.g., plain old telephone system (POTS)). Speech by the user at the server 1510 can be converted into phonemes that are then converted into multiplication 20 values that are transmitted over the transmission medium 1509 with the voice signal to facilitate the “mouthing” of words at the client 1511.

As described above, a graphical image of a human, for example, can be made to express emotions by applying a deltaset to a neutral, starting image of the human. If the expression of emotions is autonomous, the computer graphical image of the human will seem 25 more life-like. It could be concluded that humans fit into two categories or extremes: one that represents a person who is emotionally unpredictable (i.e., expresses emotions randomly), such as an infant, perhaps; and one that has preset reactions to every stimulation. According to an embodiment of the present invention, an “emotional state space” is created that includes a number of axes, each corresponding to one emotion. For example, assuming that there are 30 only two emotions, “happy” and “alert,” then at point (1.0, 0.0), the person is happy and not sleepy or excited; at point (0.0, 1.0), the person neither happy nor sad, but is excited, at point

(-1.0, -1.0), the person is sad and sleepy. Though there are many more emotions that can be expressed, a person typically will be expressing no more than one or two emotions at a time. Referring back to Fig. 9, element 937 provides input for changing the neutral image based on the expression of emotions. An example of pseudocode for the expression of emotions is  
 5 shown in Table VIII. In this pseudocode, two emotions are selected: one that is to be expressed and one that is currently fading from expression.

```
PSEUDOCODE example -----
// This is pseudocode based in part on "C".
10 // First is pseudocode for the random walk
// style of autonomous emoting, second
// is the reaction style.
// These routines determine the amount that each emotion
// in the emotion library is currently expressed
15 // in the artificial human. They do not
// actually express the emotions. One method
// of expressing the emotions is detailed above.

///////////////////////////////
20 // These variables are the basic
// output.
// emoteAmounts is an array of floats that represents
// the degree to which each emotion in the emotion
// library is currently playing on the face.
25 // emoteNum is the number of emotions in the library.
float emoteAmounts[];
int emoteNum;

/////////////////////////////
30 // These variables are the two
// emotions present at one moment.
// nextEmote & nextAmount are the
// current destination emotion &
// how much of it.
35 // lastEmote is the emotion currently
// fading away.
int nextEmote      = 0;
float nextAmount   = 0.0;
int lastEmote     = 0;
40 /////////////////////////
// This variable is the number
// of seconds it will take the
```

```
// lastEmote to fade completely.  
float decaySecs      = 3.0;  
  
5   ////////////////  
// This variable is the number  
// of seconds it will take to  
// go from the current emotion amount  
// to the next amount.  
float changeSecs     = 0.5;  
10  ////////////////  
// Routine to use a random walk to  
// navigate an emotional state-space.  
// This implementation uses only two  
15  // emotions at one time, and calls them  
// nextEmotion and lastEmotion. The  
// dynamic model is basically that of a  
// human baby, emoting at random.  
// The routine basically chooses an emotion  
20  // To go to, then increases its value while  
// decreasing the value of the previous one.  
// The input variable dt is the amount of  
// time elapsed since the last call.  
calcEmoteAmountsRandom(float dt)  
25  {  
    ////////////////  
    // These variables are probabilities  
    // of an event per second.  
    float probabilityOfNewEmote  = 0.01;  
30    float probabilityOfNewAmount = 0.2;  
  
    // Decay old emotions, go towards new.  
    DoDecayAndRamp();  
  
35    ////////////////  
    // now decide if we go to a new emotion..  
    // decide if we want to go to a new value  
    // of the current emotion without changing  
    // which emotion it is.  
40    if (unitRand()*dt <= probabilityOfNewAmount)  
        {  
            nextAmount = select a new random amount of  
emotion;  
        }  
45    // decide if we want to go to a new emotion  
    if (unitRand()*dt <= probabilityOfNewEmote)
```

```

        {
            nextEmote = a random integer >= zero and <
emoteNum;
            nextAmount = select a new random amount of
5      emotion;
        }
    }
// End of routine.
///////////////////////////////
10

///////////////////////////////
// Routine to calculate the amount
// of each emotion based on reactions
// to objects in the scene.
// This routine relies on objects with data-
// structures that contain an emotion,
// a degree of reactivity, and position.
15      CalcEmoteAmountsReact(float dt)
20      {

          // Decay old emotions, go towards new.
          DoDecayAndRamp();

          // Determine object of most interest.
25          for (i = 0; i< numberOfObjects; i++)
          {
              objectReactionLevel[i] = metric which
incorporates object's visibility, speed, speed
30          towards viewer, inherent
emotional
              reactivity (how exciting it is), and
distance to center of vision;
          }
35

          mainObject = index of largest value in
objectReactionLevel;

          // Set next Emotion & Amount.
40          nextEmotion = Object #mainObject -> reaction;
          nextAmount = objectReactionLevel[mainObject];
          }
      }
// End of routine
// Note that mainObject is also used to move the
45      artificial
          // human's eyes and head towards the object, or to

```

```

start
    // walking towards the object, and other manifestations
    // of being interested in something.
    /////////////////////////////////
5
    /////////////////////////////////
    // Routine to decay the last emotion and
    // ramp towards the next value of the
10   // new emotion.
    DoDecayAndRamp()
    {
        // Decrease value of all emotions besides current
one.
15     for (i = 0; i < emoteNum; i++)
        {
            if (i != nextEmote)
            {
                emoteAmounts[i] -= dt/decaySecs;
20                if (emoteAmounts[i] < 0.0)
                    emoteAmounts[i] 0.0;
            }
        }
        // Change value of current emotion towards
25        // next level.
        // First, calculate the direction of change.

        currAmount = emoteAmounts[lastEmote];

30        diff = nextAmount - currAmount;
        if (diff > 0.0)
            direction = 1.0;
        else if (diff < 0.0)
            direction = -1.0;
35        else
            direction = 0.0;
        // Now go in that direction at appropriate
speed.
        currAmount += dt * direction * changeSecs;
40
        // stop at ends.
        if ((direction == 1.0 AND currAmount >
nextAmount) OR
            (direction == -1.0 AND currAmount <
45   nextAmount))
            currAmount = nextAmount;

```

```
    emoteAmounts[nextAmount] = currAmount;
}
// end of decaying and ramping routine
///////////
// Utility function unitRand.
float unitRand()
```

5

10

```

{
    return a random number >= 0.0 and <= 1.0;
}

```

5

Table VIII

As seen from above, emoteAmounts [] is an array of values for the current expression of one of "emoteNum" emotions. For example, for the emotion "happy", a value is set (e.g., between -1.0 and 1.0) to indicate the current state of the graphical image (e.g., Fig. 12D shows neutralFace emoting "happy" with a value of 1.0). The nextEmote variable stores the level of the next emotion to be expressed. The lastEmote variable stores the level of the emotion that is currently being expressed, and is also fading away. The number of seconds for this emotion to fade to 0.0 is stored in the variable decaySecs. The number of seconds for the next emotion to be expressed after the current emotion amount goes to 0.0.

During the CalcEmoteAmountRandom routine, probability values for going to the next emotion (probabilityOfNewEmote) and of changing to a new amount for the current emotion (probabilityOfNewAmount) are set. Then a random number is generated, and if that number is less than the probability value, a new random amount of emotion is assigned to the variable nextAmount. A second random number is selected, and if that number is less than the probability value, a next emotion is selected from the available ones, and a random amount is assigned to the nextAmount variable.

During the routine CalcEmoteAmountsReact, the objects that are around the graphic image of the person are analyzed to determine which object is of most interest (e.g., by assigning weighted values based on the object's visibility, speed, speed towards the graphical image of the person, the inherent emotional reactivity of the object, and its distance to center of vision for the graphic image of the person). Each object has a data structure that includes a predefined emotion, a degree of reactivity and position. For example, a gun object, would elicit a "fear" emotion with a high degree of reactivity depending on how close it is (i.e., distance) to the person. Accordingly, based on the object of most interest (and the relationship between the person and the object), a nextEmotion is selected and a nextAmount is selected based on the object and the random numbers referenced above determine whether that next Emotion is to be expressed by the human image. Using the routines of Table VIII, the human

image expresses emotions that are more lifelike in that they are somewhat random, yet can occur in response to specific stimuli.

Referring back to Fig. 8, an input device 830 is provided for the input of data for the creation of graphic images to be output to display 820. The input device 830 can be a variety 5 of components including a video camera, a magnetic tracker monitor, etc. In one such system, selected points are tracked on a person's face. These devices output a stream of information that are commensurate with the coordinates of a number of select locations on a person's face as they move (see element 939 in Fig. 9). For example, 6 locations around the mouth, one on each eyelid, one on each eyebrow, one on each cheek, can all be tracked and output to the 10 computer system of Fig. 8.

The method of face tracking according to an embodiment of the present invention, a neutral three-dimensional model of a person is created as described above. A test subject (e.g., a person) is used having a set of markers on his/her face (as described above). For each marker, three three-dimensional model faces are created, one for each 3D axis (e.g., the x, y 15 and z axes). Each of these models is the same as the neutral model except that the specific marker is moved a known distance (e.g. one inch or other unit) along one of the axes. Thus, for each marker, there is a contorted version of the neutral image where the marker is moved one unit only along the x-axis; a second image where the marker is moved along one unit only along the y-axis; and a third image where the marker is moved along one unit only along 20 the z-axis. Deltatasets are then created between the neutral image and each of the three contorted versions for each marker.

With the deltatasets created, the input stream of marker positions are received from the input device 830. The neutral image is then modified with the appropriate deltaset(s) rather than directly with the input positions. If marker data is only in two dimensions, then only two 25 corresponding distorted models are needed (and only two deltatasets are created for that marker). Movement of one marker can influence the movement of other points in the neutral model (to mimic real-life or as desired by the user). Also, the movement of a marker in one axis may distort the model in more than one axis (e.g., movement of the marker at the left eyebrow in a vertical direction may have vertical and horizontal effects on the model). An 30 example of pseudocode for implementing the input of marker positions is shown in Table IX.

```

// This pseudocode is based in part on "C".
// It takes as input:
//   * An array of vector data representing the spacial
//     displacement of a set of facial markers.
5 //   * An array of DeltaSets setup as described above,
//     with numDimensions DeltaSets for each marker.
//   * a 3D model of a "neutral face", as described above.
// It outputs:
//   * a 3D model of a face which mimics the motion of the
10 // actual face with which the markers are associated.

    int      numMarkers;
    int      numDimensions;
    float    markerDisplacements [numMarkers]
15 [numDimensions];
        DeltaSet  markerDeltaSets [numMarkers] [numDimensions];
        Shape     neutralFace;
        Shape     outputFace;
        // numMarkers is the number of discrete locations being
20 // tracked on the source face.  Typically 6-14, but
        // under no limitations.
        // numDimensions is the number of dimensions reported by
        // the markerDisplacements array.
        // markerDisplacements is an array of vectors with one
25 vector
        // for each marker on the source face.  These values
        should
        // be updated once per frame.
        // marker DeltaSets is a 2D array of DeltaSets of size
30 // numMarkers x numDimensions.
        // neutralFace is the original, undistorted 3D face
        model.
        // outputFace is a 3D model that will mimic the source
        face.
35
////////// // The main animation loop.  Runs once per frame.
        MainAnimationLoop ()
{
40        outputFace = copy of neutralFace;

        // Loop over each marker and each reported
        dimension.
        for (m = 0; m < numMarkers; m++)
45        {
            for (d = 0; d < numDimensions; d++)

```

```
        {
            deltaSet_Apply (markerDeltaSets [m] [d] ,
                            & outputFace,
                            MarkerDisplacements
5           [m] [d]) ;
        }
    }
//  
10 // End of main animation loop.  
  
//////////
```

15

## Table IX

As seen from the above, the neutral face image that is input is modified with the created deltaset to mimic the resultant movements in the face caused by physically moving the attached markers. Without distortion, neutralFace is the original 3D face model and  
20 outputFace is a 3D model that mimics the movement of the subject's face. During the main animation loop, which can run once per frame, each marker is analyzed for its position. The resulting displacement of the marker is then applied to the outputFace (which starts as a copy of neutralFace) through the use of the Deltaset\_apply routine discussed above and the deltaset that have been previously created.

What is claimed is:

1. A method of transforming a graphical image, comprising:
  - (a) identifying a first region of a first graphical image; and
  - (b) modifying said first region of said first graphical image based on a first set of predetermined modification data.
2. The method of claim 1 further comprising after step (a):
  - (a1) retrieving said first set of modification data from a database, said database storing a plurality of sets of modification data.
3. The method of claim 1 further comprising:
  - (c) identifying a first region of a second graphical image; and
  - (d) modifying said at least first region of said second graphical image based on said first set of predetermined modification data.
4. The method of claim 1 further comprising:
  - (c) further modifying said first region of said first graphical image based on a second set of modification data.
5. The method of claim 1 further comprising:
  - (c) identifying a second region of said first graphical image; and
  - (d) modifying said second region of said first graphical image as modified in step (b) based on a second set of modification data.
6. A method of transforming a starting graphical image to a destination graphical image, comprising:
  - (a) generating a first image;
  - (b) generating a second image;
  - (c) determining a set of difference values between said first image and said second image; and
  - (d) storing said difference values in a memory device.

7. A method of transforming a starting graphical image to a destination graphical image, comprising:
  - (a) generating a first image comprising a first number of vertices, each vertex having a spatial location;
  - (b) generating a second image comprising a number of vertices equal to said first number of vertices;
  - (c) determining a difference between a first one of said vertices of said first image with a corresponding vertex of said second image; and
  - (d) storing said difference in a memory device.
8. A method of modifying a starting graphical image comprising:
  - (a) identifying at least a first region of said starting graphical image;
  - (b) modifying said first region of said starting graphical image based on first and second sets of modification data.
9. The method of claim 8, further comprising before step (b):
  - (a1) adding together said first and second sets of modification data.
10. A communication system, comprising:  
a client including a computer system and a display;  
a server including at least a computer system, said server capable of transmitting a first image and at least a first set of modification data to said client; and  
said server is capable of modifying said first image by transmitting a multiplier value to said client, such that said first image is modified based on said first set of modification data multiplied by said multiplier value.
11. A method of modifying a starting graphical image comprising:
  - (a) generating a starting image having a plurality of characteristics;
  - (b) storing in a library a plurality of sets of modification data;
  - (c) retrieving a first set of modification data from said library; and
  - (d) modifying said starting image based on said first set of modification data, such

that a first characteristic of said starting image is changed.

12. The method of claim 11 further comprising after step (c):

(c1) multiplying said first set of modification data by a multiplier value;

and further wherein in step (d), said starting image is modified based on said first set of modification data multiplied by said multiplier value.

13. The method of claim 12 wherein said multiplier value is input by a relative position of a graphical slider.

14. The method of claim 11, further comprising repeating steps (c) and (d) for each change to one of said characteristics of said starting image.

15. The method of claim 12, further comprising repeating steps (c), (c1), and (d) for each change to one of said characteristics of said starting image.

16. The method of claim 14 wherein said multiplier value is input by a relative position of a graphical slider.

17. A method of modifying a starting graphical image comprising:

(a) generating a first starting image having a plurality of characteristics;

(b) storing in a library a plurality of sets of modification data;

(c) retrieving first and second sets of modification data from said library; and

(d) modifying said starting image based on said first and second sets of modification data, such that first and second characteristics of said starting image are changed, respectively.

18. The method of claim 17 further comprising after step (c):

(c1) multiplying said first set of modification data by a first multiplier value; and

(c2) multiplying said second set of modification data by a second multiplier value; and further wherein in step (d), said starting image is modified based on said first and second

sets of modification data multiplied by said first and second multiplier values, respectively.

19. The method of claim 18 wherein said first and second multiplier values are input by relative positions of first and second graphical sliders, respectively.

20. The method of claim 18 further comprising repeating steps (c), (c1), (c2), and (d) for each change to said first and second characteristics of said starting image.

21. The method of claim 17, further comprising:

(e) generating a second starting image having a plurality of characteristics;  
(f) retrieving said first and second sets of modification data from said library; and  
(g) modifying said second starting image based on said first and second sets of modification data, such that first and second characteristics of said second starting image are changed, respectively.

22. The method of claim 21, further comprising after step (f):

(f1) multiplying said first set of modification data by a first multiplier value; and  
(f2) multiplying said second set of modification data by a second multiplier value.

23. A method of modifying a graphical image comprising:

(a) generating a starting graphical image of a character, said character capable of expressing a plurality of emotions;  
(b) retrieving a first set of modification data from a library, each set of modification data stored in said library being related to an expression of an emotion;  
(c) multiplying said first set of modification data by a first multiplication value; and  
(d) modifying said starting graphical image based on said first set of modification data multiplied by said multiplier value.

24. The method of claim 23, further comprising:

(e) generating a first random number;

- (f) comparing said first random number with a first probability value;
- (g) selecting a new amount for said first multiplication value based on said comparison in step (f); and
- (h) repeating steps (c) and (d).

25. The method claim 24, further comprising:

- (e) generating a second random number;
- (f) comparing said second random number with a second probability value;
- (g) selecting a second set of modification data based on said comparison in step (f); and
- (h) multiplying said second set of modification data by a second multiplication value; and
- (i) modifying said starting graphical image based on a said second set of modification data multiplied by said second multiplication value.

26. The method claim 24, further comprising:

- (i) generating a second random number;
- (j) comparing said second random number with a second probability value;
- (k) selecting a second set of modification data based on said comparison in step (j);
- (l) multiplying said second set of modification data by a second multiplication value; and
- (m) modifying said starting graphical image based on a said second set of modification data multiplied by said second multiplication value.

27. A method of modifying a graphical image comprising:

- (a) generating a starting graphical image of a character, said character capable of expressing a plurality of emotions;
- (b) generating at least a first graphical object located proximately to said starting graphical image, said graphical object including a data structure defining a first emotion to be expressed by said starting graphical image;

- (c) determining a relationship between said starting graphical image and said graphical object;
- (d) retrieving a first set of modification data from a library, said modification data being related to an expression of a said first emotion;
- (e) generating a multiplication value based on said relationship between said graphical object and said starting graphical image;
- (f) multiplying said first set of modification data by said multiplication value; and
- (g) modifying said starting graphical image based on said first set of modification data.

28. The method of claim 27 wherein said relationship is a distance between said graphical object and said starting graphical image.

1/15

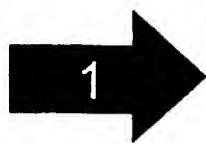
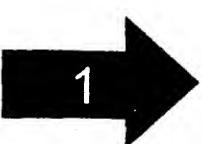
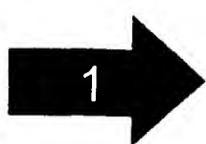


FIG. 1

SUBSTITUTE SHEET (RULE 26)

BEST AVAILABLE COPY

2/15

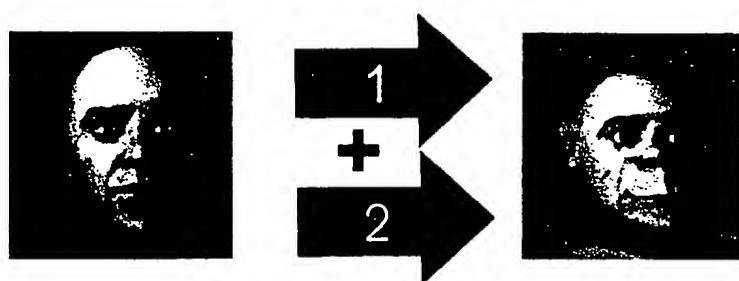
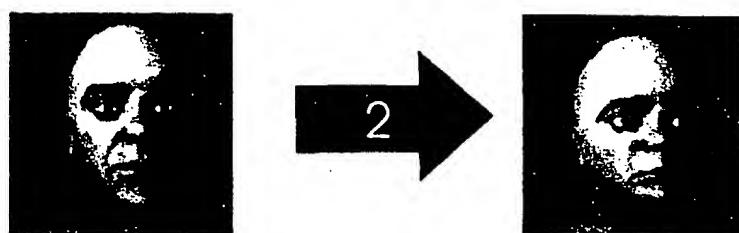
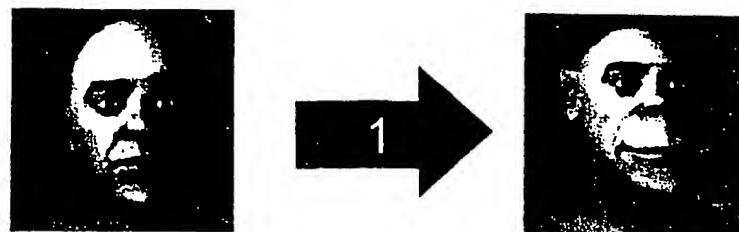


FIG. 2

SUBSTITUTE SHEET (RULE 26)

BEST AVAILABLE COPY

3/15



4-VISEME

"P"

"O"

"N"

"Y"

---

3-TIME →

FIG. 3

SUBSTITUTE SHEET (RULE 26)

BEST AVAILABLE COPY

4/15

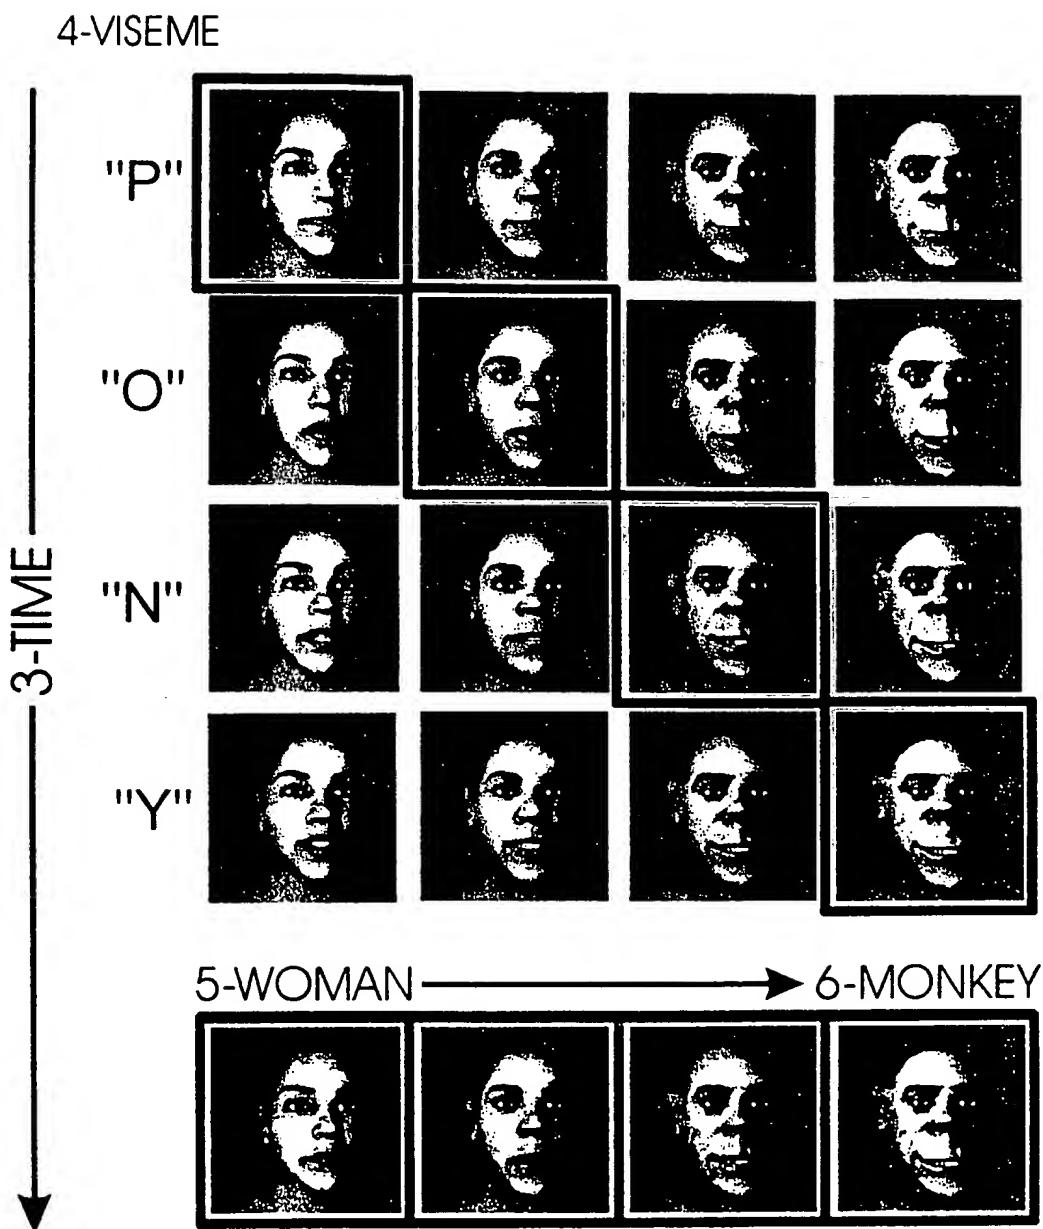


FIG. 4  
SUBSTITUTE SHEET (RULE 26)

5/15

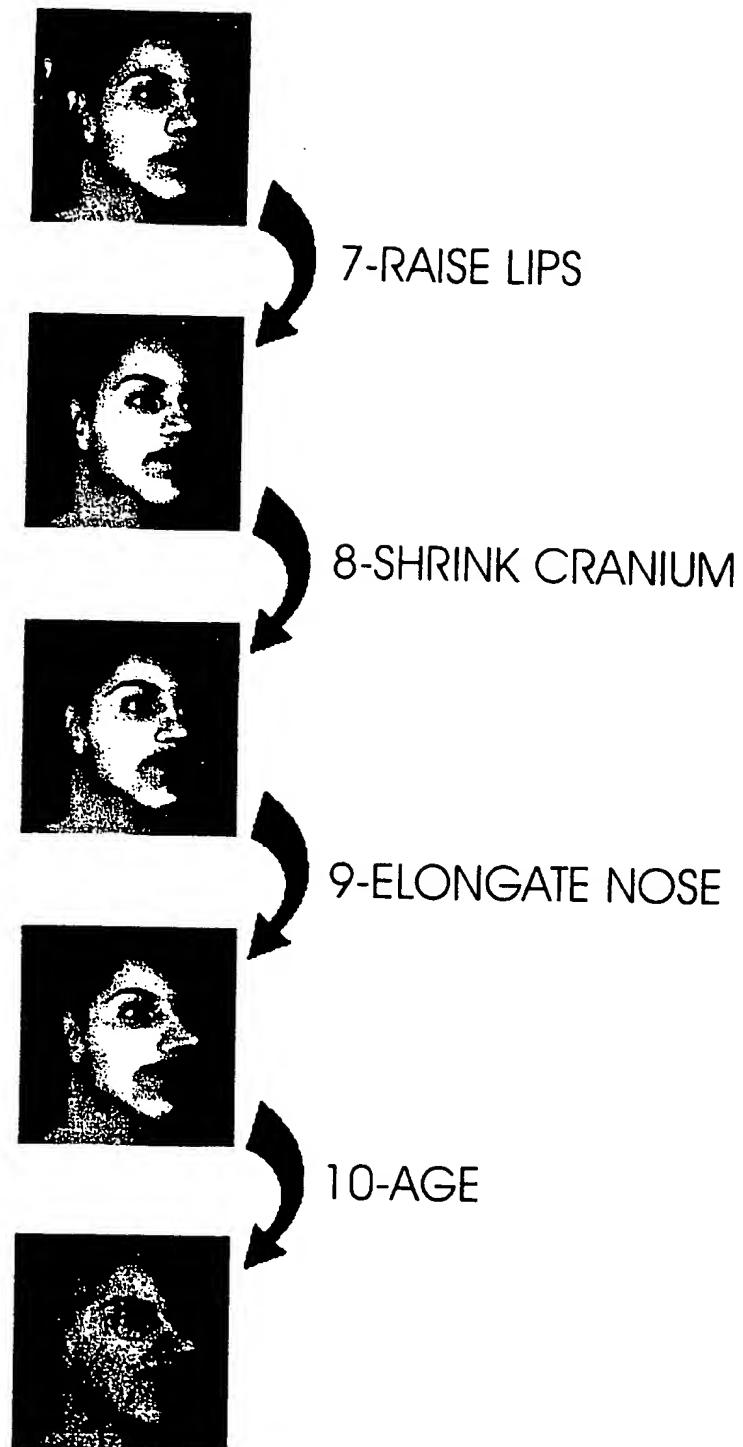


FIG. 5  
SUBSTITUTE SHEET (RULE 26)

6/15

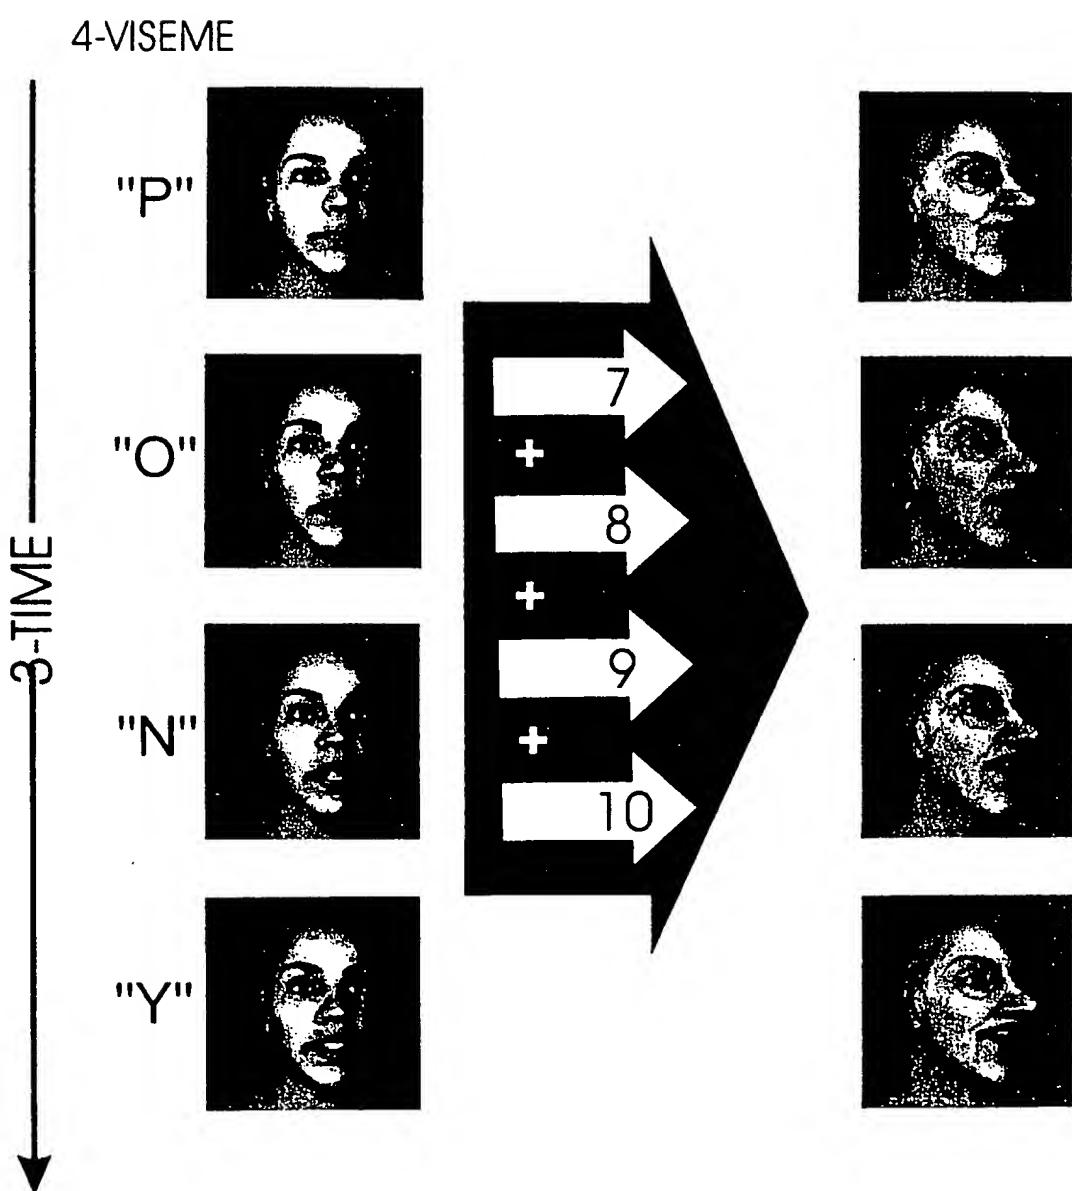


FIG. 6  
SUBSTITUTE SHEET (RULE 26)

7/15

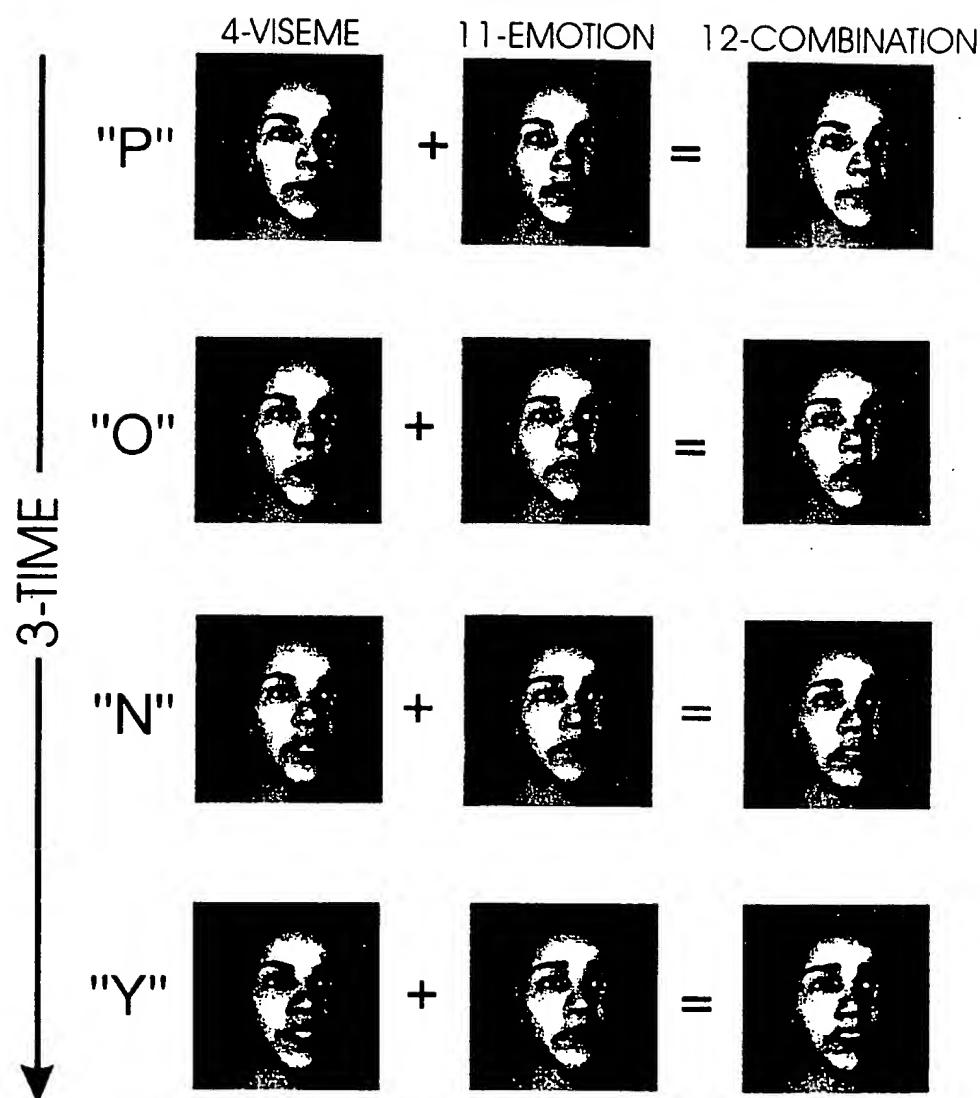


FIG. 7  
SUBSTITUTE SHEET (RULE 26).

8/15

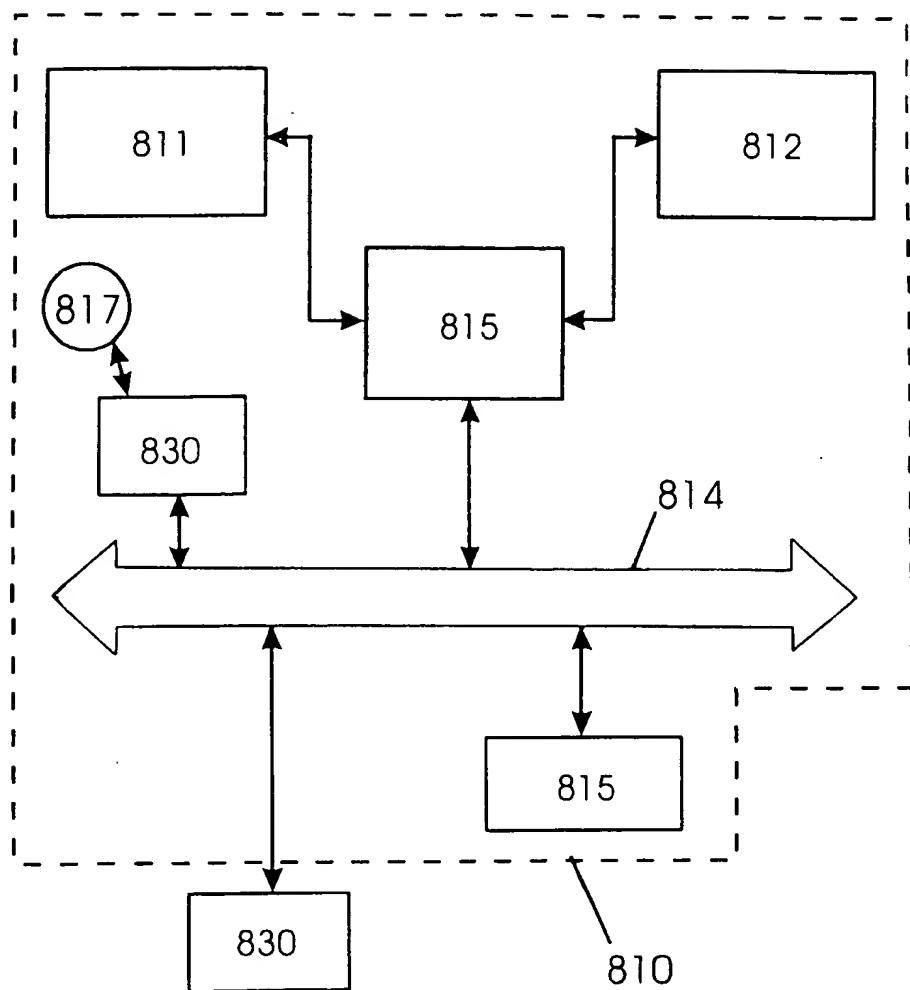
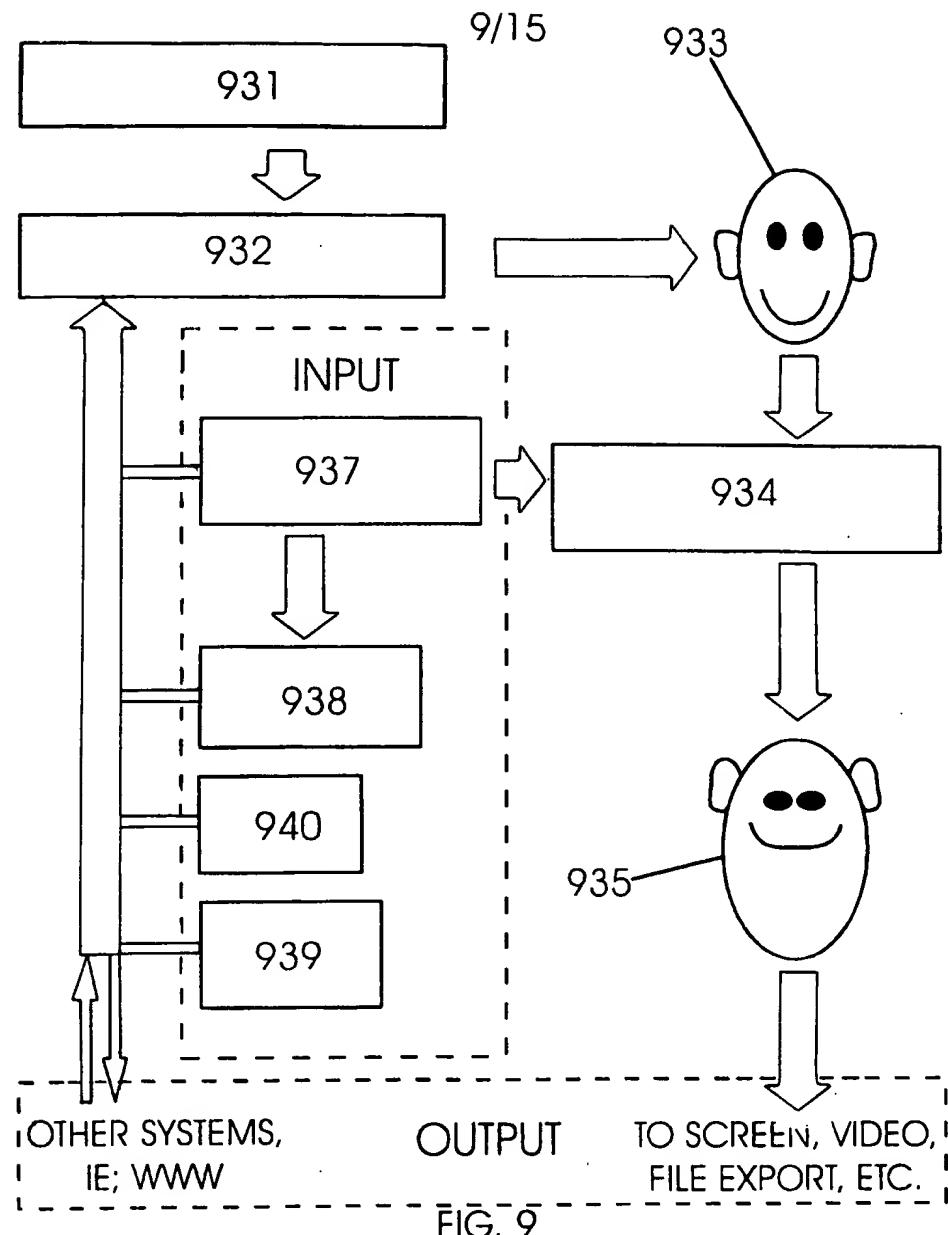


FIG. 8



10/15

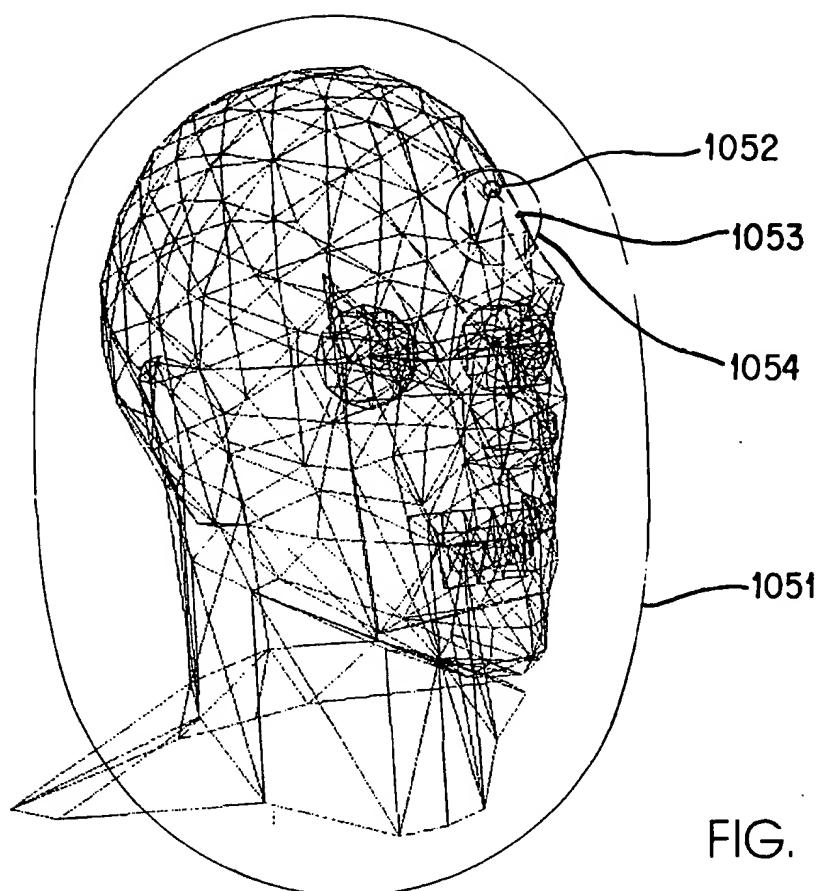


FIG. 10A

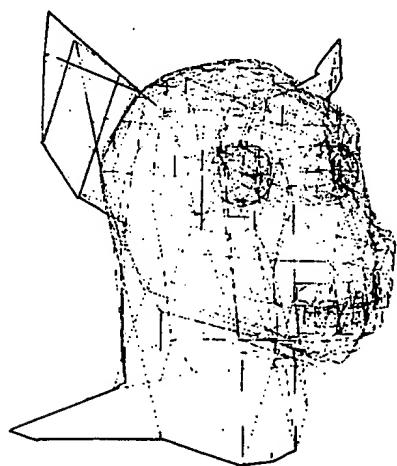


FIG. 10B

SUBSTITUTE SHEET (RULE 26)

11/15

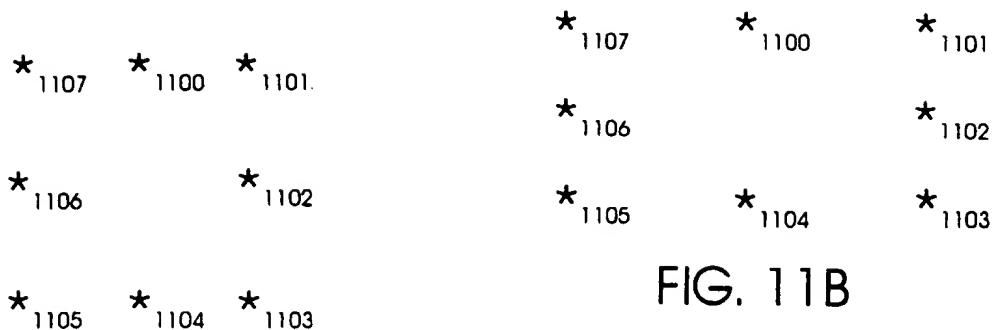


FIG. 11A

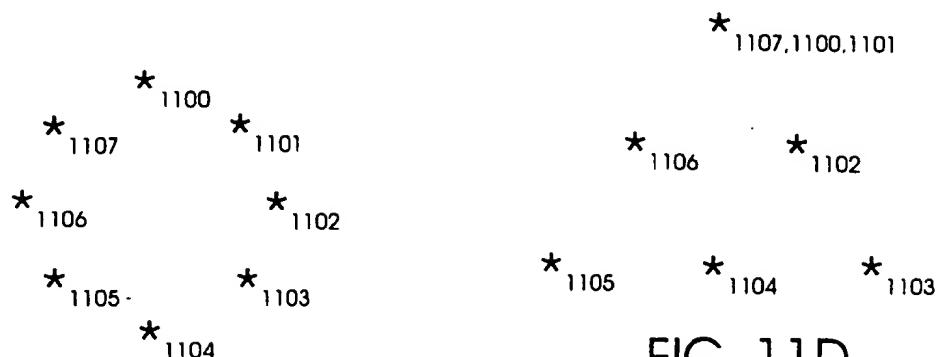


FIG. 11D

FIG. 11C

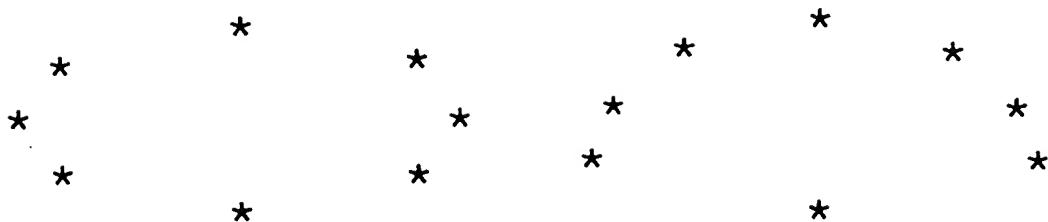


FIG. 11F

FIG. 11G

12/15



FIG. 12A



FIG. 12B



FIG. 12C



FIG. 12D

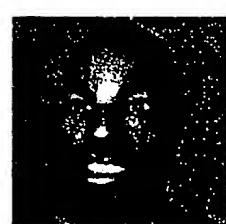


FIG. 12E



FIG. 12F



FIG. 12G

13/15

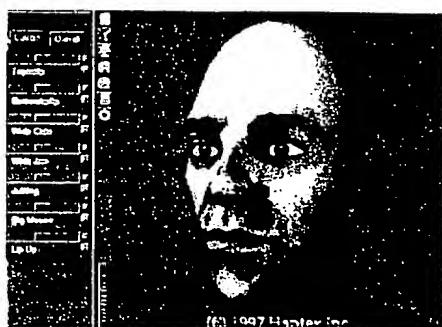


FIG. 13A



FIG. 13B

1301

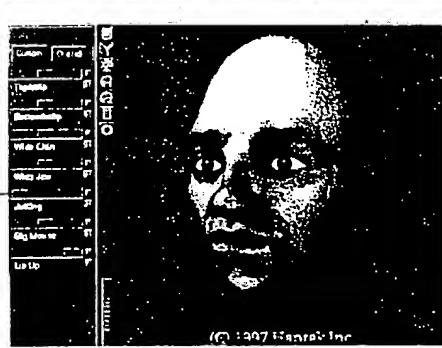


FIG. 13C

1302

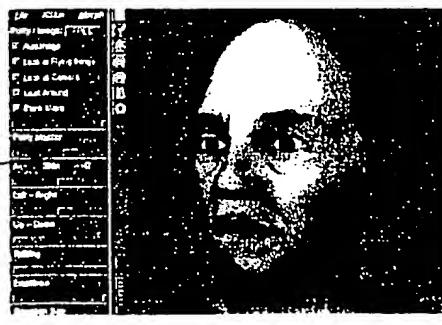


FIG. 13D

1303

14/15



FIG. 14A



FIG. 14B



FIG. 14C



FIG. 14D

SUBSTITUTE SHEET (RULE 26)

BEST AVAILABLE COPY

15/15

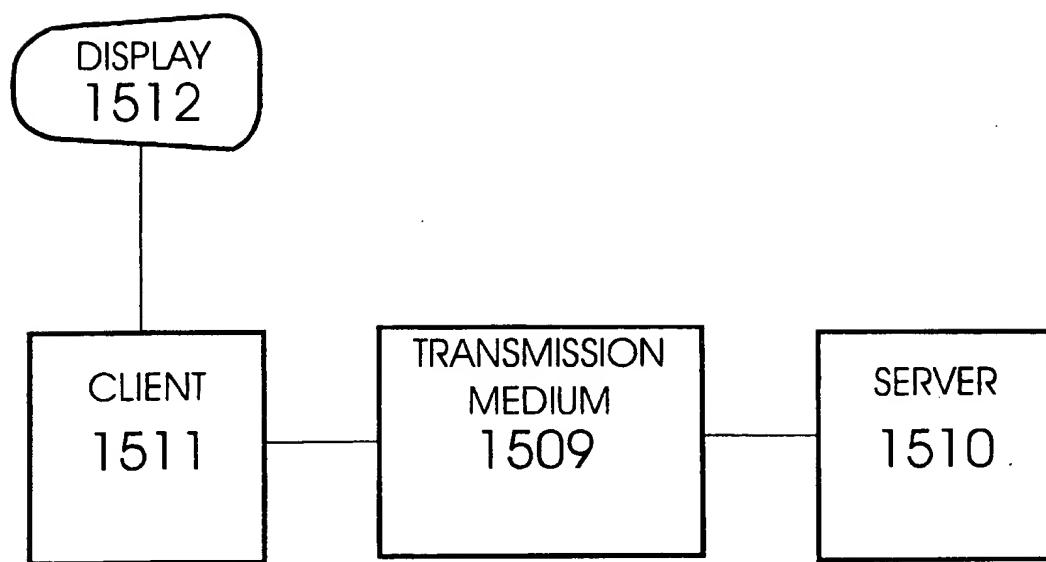


FIG. 15

**This Page is Inserted by IFW Indexing and Scanning  
Operations and is not part of the Official Record**

## **BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** \_\_\_\_\_

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.**